

Alexandria: Extensible Framework for Rapid Exploration of Social Media[†]

Fenno F. Heath III, Richard Hull, Elham Khabiri,
Matthew Riemer, Noi Sukaviriya, and Roman Vaculín
IBM Research

Yorktown Heights, New York, USA

Email: {theath, hull, ekhabiri, mdriemer, noi, vaculin}@us.ibm.com

Abstract—The Alexandria system under development at IBM Research provides an extensible framework and platform for supporting a variety of big-data analytics and visualizations. The system is currently focused on enabling rapid exploration of text-based social media data. The system provides tools to help with constructing “domain models” (i.e., families of keywords and extractors to enable focus on tweets and other social media documents relevant to a project), to rapidly extract and segment the relevant social media and its authors, to apply further analytics (such as finding trends and anomalous terms), and visualizing the results. The system architecture is centered around a variety of REST-based service APIs to enable flexible orchestration of the system capabilities; these are especially useful to support knowledge-worker driven iterative exploration of social phenomena. The architecture also enables rapid integration of Alexandria capabilities with other social media analytics system, as has been demonstrated through an integration with IBM Research’s SystemG. This paper describes a prototypical usage scenario for Alexandria, along with the architecture and key underlying analytics.

I. INTRODUCTION

Twitter, Instagram, forums, blogs, on-line debates, and many other forms of social media have become the outlets for people to freely and frequently express ideas. Indeed, many research papers have explored social media usage in many application areas. Research has ranged from using social science techniques to find indicators of phenomena such as increased health risks, to studies on optimization of hugely scaled analytics computations, to usability of analytics visualizations. However, there has been little work on how to bring together the myriad of analytics capabilities to support knowledgeable business analysts in rapid, collaborative, and iterative exploration and analysis of large data sets. This requires a combination of several aspects, including integration of numerous analytics tools, efficient and scalable data and processing management, a unified approach for data and results visualization, and strong support for on-going knowledge-worker driven activity to uncover and focus in on particular areas of interest. The paper describes the Alexandria system, currently under development at IBM Research, which supports these several aspects. The system is currently focused on the early stages of the overall analytics lifecycle, namely, on enabling rapid, iterative exploration and visualization of social media data in connection with a given domain (e.g., consumption habits for beverages, the growth of the market for vegan foods, or political opinions about an upcoming election). The system has been designed to

support rich extensibility, and has already been integrated with a complimentary system at IBM.

Figure 1 shows the two main parts of current Alexandria processing, namely Background Processing and Iterative Exploration. The Background Processing includes primarily (a) various analytics on background text corpora that support several functionalities, including similar term generation, parts-of-speech and collocation analytics, and term-frequency-inverse-document-frequency (TF-IDF) analytics; and (b) ingestion and indexing of social media data (currently from Twitter) to enable main-memory access speeds against both text and structured document attributes. (Although not shown in the figure, there is also background analytics to compute selected author profile attributes, e.g., geographic location, family aspects, interests). Iterative Exploration enables users to build a number of related *Projects* as part of an investigation of some domain of interest. Each Project includes (i) the creation of a targeted *domain model* used to focus on families of tweets and authors relevant to the investigation, (ii) application of a variety of analytics against the selected tweets and their authors, and (iii) several interactive visualizations of the resulting analytics. At the beginning of an investigation there are typically several *experimental* Projects, used by individuals or small collaborating groups. Over time some Projects may be *published* with more stability for broader usage.

Alexandria advances the state of the art of social media analytics in two fundamental ways (see also Section VIII). First, the system brings together several text analytics tools to provide a broad-based environment to rapidly create domain models. This contrasts with research that has focused on perfecting such tools in isolation. Second, Alexandria applies data-centric and other design principles to provide a working platform that supports ad hoc, iterative, and collaborative exploration of social media data. As such, the system extends upon themes presented in [19], [8], and develops them in the context of social media applications.

Section II highlights the key goals for Alexandria, including both longer- and shorter-term ones. Section III describes a prototypical usage scenario for the system, and illustrates its key functionalities. Section IV highlights key aspects of the system architecture, and describes how the design choices support the key goals. Section V describes key technology underpinnings for the domain scoping capability, and Section VI does the same for the currently supported analytics. Section VII describes the data-centric approach taken for managing exploratory Projects to enable rich flexibility. Section VIII describes related work, and Section IX discusses future directions.

[†] This material is based upon work supported by the U.S. Defense Advanced Research Projects Agency (DARPA) under Agreement Number W911NF-12-C-0028. The content of the information in this document does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation here on.

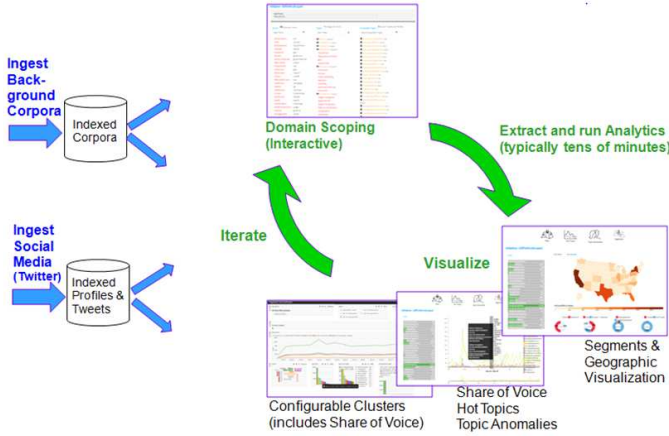


Figure 1. Alexandria supports iterative exploration of social media, and includes background text analytics processing. (See also Figure 9.)

II. SYSTEM GOALS

This section outlines the primary long- and shorter-term goals that have motivated the design of the Alexandria framework and system.

The longer-term goals are as follows:

LG1: Extensible platform to support business users with numerous styles of analytics. This contrasts significantly with most previous works, that are focused primarily on scalable performance, support for targeted application areas, or support primarily for data scientists. Alexandria is focused on providing a layer above all of these, to enable business users to more effectively use analytics, both to find actionable insights, and also to incorporate them into on-going business processes.

LG2: Support analytics process lifecycle, from exploration to prescription. As discussed in [8], there are several stages in the lifecycle of analytics usage, ranging from initial exploration, to refinement and hardening, to incorporation into already existing business processes for continuing value add, to expanding the application to additional aspects of a business. While the CRISP-DM method [21] addresses several elements of the lifecycle, the method and associated tools are geared towards data scientists rather than business users. In contrast, a goal of Alexandria is to provide business users with substantial exploration capabilities, and also support the evolution of analytics approaches from exploration to production usage. Of course data scientists will still play a very key role, and the Alexandria platform should enable graceful incorporation of new algorithms as they become available from the data scientists.

LG3: Support for a collaborative production environment. Analytics is no longer the realm of a small team of data scientists working largely in isolation. Rather, it is increasingly performed by a multi-disciplinary team that is in parallel digging more deeply into the data, finding ways to add business value by integrating analytics insights into existing business processes, and finding ways to make the usage of the insights production grade.

LG4: Scalable, e.g., work with billions of tweets and forum comments. The Alexandria system should be able to work with state-of-the-art systems such as SPARK and TITAN, and more generally with Hadoop-based and other

distributed data processing systems, to enable rapid turnaround on large analytics processes. Similarly, the system should support main-memory indexing systems such as Elastic Search or LUCENE/SOLR to enable split-second access from very large data sets, including text-based searches.

As a way to get started with the longer-term goals, the initial version of Alexandria has focused more narrowly on (a) Social Media analytics, and on (b) the exploration and initial visualization phases of the overall analytics process. The key shorter-term goals include the following:

SG1: Enable users to begin their exploration of a new topic domain within a matter of hours.

SG2: In particular, enable non-experts to quickly create a domain model (i.e., keywords and extractors) that enables a focus on Tweets and other social media that are relevant to a given topic.

SG3: Provide a variety of different analytics-produced views of the data, to permit different styles of data and results examination

SG4: Support iterative exploration based on info learned so far, including management of meta-data about raw and derived data sets

SG5: Minimize processing time through to enable as much interactivity as possible, by using main-memory indexes, parallel processing, avoiding data transfers, etc.

SG6: Enable easy and fast orchestration of capabilities, including rapid creation of variations on the domain model and the analytics processing. This includes the automation of processing steps and the defaulting of configuration parameters wherever possible.

III. USING THE SYSTEM

This section illustrates the main capabilities currently supported in Alexandria through an extended example.

To extract relevant documents from social media, one needs to gather documents that mentioned terms, expressions, or opinions pertaining to the area one wants to explore. Alexandria provides tools that support both laymen and experts in finding terms that cover the space of interest, and also terms that can drill more deeply into that space.

We will explore a subject around vaccination as an example for this paper. Suppose that the government would like to encourage people to take vaccination, but wonder what peoples opinions may be around vaccination. The exploration starts with creating a Project with a few seed terms, namely ‘vaccination’, ‘flu’ and ‘measles’. Based on these terms, we asked Alexandria to generate a family of relevant collocated terms in an effort to bound the scope. These terms may be manually edited, to reach the terms listed in Figure III. Here, the black terms were generated automatically, red were added by hand, and gray with strike out were auto-generated but deleted by hand.

In some cases the auto-generated terms will help the user learn more about the domain of interest. In this example, Dr. Anne arises, and a Google search reveals that Dr. Anne Schuchat is the Director of the U.S. Center for Disease Control [20], so her name was left in the list. Similarly, Dr. Gil remains because he is mentioned in a news article [11] concerning a measles outbreak at Disneyland.

While the scoping step is supposed to extend our vocabulary to cover various areas of the topics, some terms

vaccination, measles, immunization, CDC, autism, vaccines, vaccinations, vaccine, diseases, whooping, cough, accessed, flu, children, preventable, leaks, immunity, vaccinating, immunizations, adults, recommended, schedule, rates, Pinellas, pertussis, infectious, parents, California, vaccinated, resurgence, harmful, cervical, moms, cases, rate, mom, risk, hepatitis, mumps, mmr, vaccination rates, National Center, Immunization Respiratory, Mom, Department Public, March-Dimes, United States, exemption%, vulnerable outbreaks, infectious diseases, people children, Dr. Anne, slight dip, contagious key, surgeon general, Landsman Aeon, California parents, throat penile, diseases vaccine-preventable, personal belief, public acceptance, resurgence unambiguous, flu obvious, Learn fact, rhetoric recent, Schuchat assistant, country cases, measles ongoing, vaccination community, year measles, preventable kinds, t great, schools California, January August, Dr. Gil, greatest U.S., months milestone, dirt global, percentage kindergartens, Diseases resurgence, advocates small, small harmful, officials insufficient, director CDC, U.S. case, latest figures, LAT challenge, movement movement, Centers Disease, diseases measles

Figure 2. List of relevant terms collocated with the seed terms, after manual edits

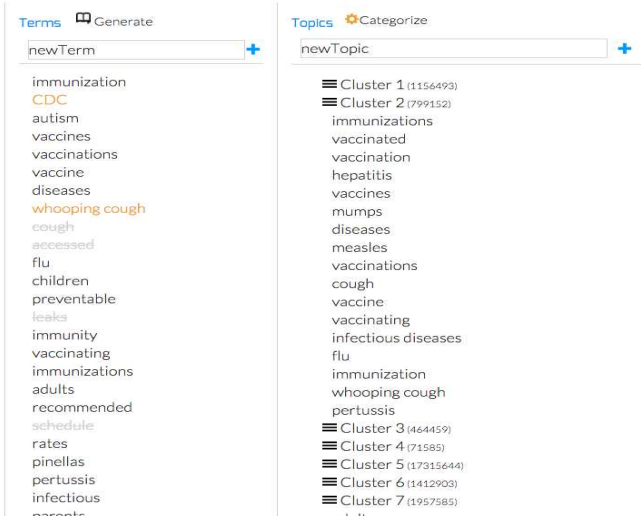


Figure 3. Alexandria interface for domain scoping: After automatic term clustering

appear to be rather similar. For example, many variations of vaccination are included in the list. We know that if a tweet mentioned one of these terms, it is likely to have something to do with vaccination. Alexandria supports automatically clustering similar terms into groups called topics. Each topic is used to provide a list that, if a tweet mentions one or more of the terms in the topic, we can conclude that the tweet has mentioning of this topic.

Figure 3 above shows a snippet from the actual Alexandria page where the terms are listed vertically in the first column and the second column shows the clusters suggested by Alexandria. Note that these clusters are generically named Cluster 1, Cluster 2, and so forth. In the figure Cluster 2 is “open”, to show terms Alexandria placed in it. It appears many vaccination and diseases that can be prevented by vaccination are included. In Section V, we detail the analytics we are doing behind the scenes for topic clustering.

The numbers to the right of each topic indicate the number of tweets in which at least one term from the topic is found. One can use this number to gauge how widespread the topic is. Bear in mind that something general such as Disbelieved can be about any subject, hence the large number of over a million tweets, and not necessarily about vaccination. These numbers are obtained within seconds through accesses to a

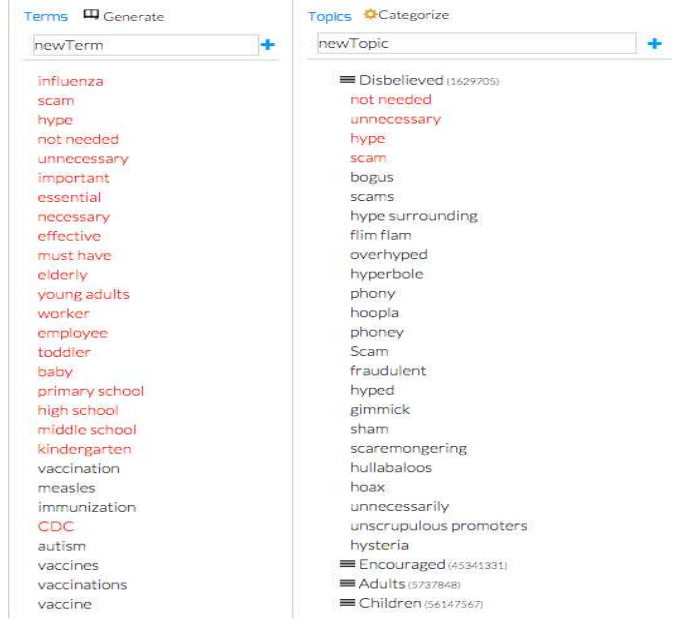


Figure 4. Topics after adding terms in the “Disbelieved” topic

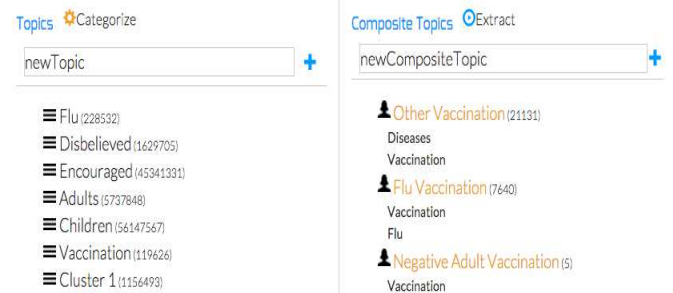


Figure 5. Topics after adding terms in the “Disbelieved” topic

SOLR index holding all of the tweet information.

Figure 5 illustrates the state of the system after a few steps. First, Alexandria supports renaming of the clusters, and moving them around in the middle column. Second, there is an automated “similar term generation” service for adding depth to an existing topic. In the figure, the red terms in Disbelieved were inserted by hand, and the terms in black below that were generated automatically to add depth.

The third phase of scoping is the building of the actual extractors (or queries) for selecting tweets of interest. This is accomplished by creating composite topics, which are based on Boolean combinations of the topics. Figure 5 shows several composite topics, some of which are “open” to expose the topics that are used to form them. (At present the UI supports only conjunctions of topics, but the underlying engine supports arbitrary combinations.) For example, a composite topic Support Flu Vaccination we combine Flu, Vaccination and Encouraged topics to form a search statement of find any tweets that mentioned at least one (or more) of the terms in Flu and one (or more) of the terms in Vaccination and one (or more) of the terms in Encouraged. (A further refinement would be to exclude tweets that include a negating term such as “not”.)

Once the set of composite topics has been specified, it is

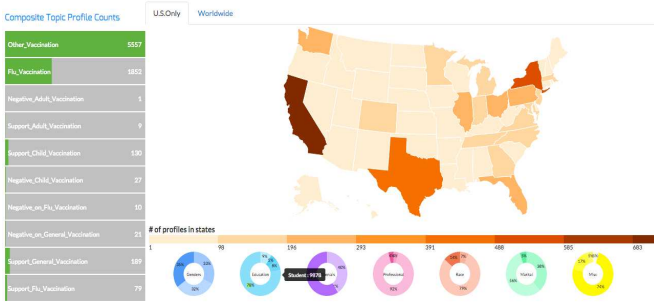


Figure 6. Interactive view for exploring the demographic distribution of tweet authors who are negative about flu vaccination

time to perform some data extractions, re-structurings, and indexing to support various analytics. Upon request, Alexandria extracts tweets with topics matching the composite topic combinations, annotates each tweet accordingly, and then launches multiple analytics activities on these tweets. One of the activities was extracting the author profiles of these tweets and aggregate attributes among these profiles. We will detail this work on in Section VI on Analytic View.

We now describe some of the visualizations used to show the analytics results associated with a Project. In one direction, Alexandria infers profile attributes of Twitter authors through background analysis of 100's tweets per author. Information such as education, gender, ethnicity, location of residence is inferred based on evidence of words found in tweets. Figure 6 shows how the demographic distribution of tweet authors of composite topics in the U.S. On the left, it shows the numbers of authors for various composite topics. On the map, states with darker colors mean higher numbers of authors reside in those states. Mousing over a state (not shown) would give more details of these authors. The colored donuts below the map show percentage of various characteristics of those located in the U.S. for example, male, female or unknown for gender. Mousing over a portion of a donut shows the value of the characteristics and the number of profiles. For example, in the figure we show that 5898 tweet authors of all topics combined are students.

Figure 7 illustrates another analytic view in which Alexandria shows share of voices, i.e., comparison of tweet volumes of the composite topics over time. In this paper we are working on tweets from January to June of 2014. Notice the higher volumes among the topics Flu vaccination and Other Vaccination in Figure 7, with a peak around mid-May for Other vaccination topic. One may wonder what happened during that week. In this view we can click on the graph to explore the frequently mentioned terms or anomalous terms mentioned in that week. Figure 8 shows snippets of two images captured to highlight the two types of terms.

Specifically for Figure 8, we selected the Flu Vaccination topic on the left to narrow the visualization down to just this topic, hence the presence of only one line graph in the two snippets. This line represents the volume over time of tweets that match the Flu Vaccination extractor. For this topic, there seems to be a peak around the second week in January. The snippet on the left of the figure shows frequently mentioned terms in the week while the snippet on the right shows terms that are considered anomalous in that week. We moused over the term swine flu outbreak which was mentioned 19 times, hence showing up high in the word list. However,




Figure 7. Share of voices of tweets from different composite topics

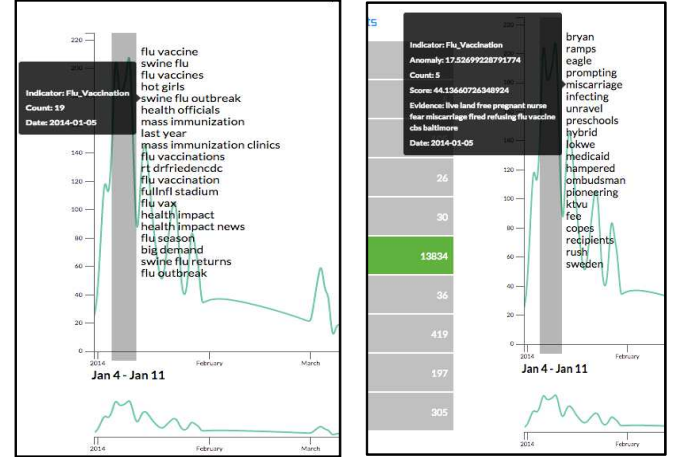


Figure 8. Exploring frequently mentioned and anomalous terms for the composite topic Flu Vaccination in the week of Jan4 to Jan 11. The text boxes show contextual data of the term they point to. The box on the right also shows partial context of the tweets where the term was extracted from.

this term is not considered anomalous, indicating that this term also shows up fairly often in other weeks. However, the term miscarriage is anomalous. Mousing over it reveals some evidence of the news about a nurse refusing to get vaccinated and subsequently being fired from a hospital.

There are other views that one can use to explore analytics results of social media insights, including some that leverage the configurable Banana visualization tool [1].

IV. SYSTEM ARCHITECTURE

The Alexandria architecture will be described from three perspectives: (a) the overall processing flow (see Figure 9), (b) the families of REST APIs supported (Figure 10), and (c) the key systems components. These descriptions will include discussion of how the architectural choices support the long-term and short-term goals.

Overall, the current Alexandria architecture flow shown in Figure 9 expands on Figure 1, and is focused on supporting rapid exploration, analytics processing, and visualization of Twitter data in a collaborative environment, that is, on parts of goals LG2, LG3 and LG4, and all of goals SG1 through SG6. There are two forms of background processing. One is to ingest and index the Tweets, and also includes author-by author processing of tweets to extract demographic attributes, such as gender, geographic location, and one to ingest, process, and index background text corpora. (This demographic processing uses the IBM Research SMARC

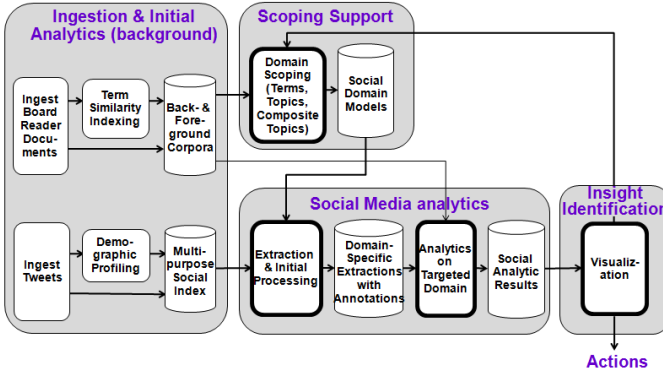


Figure 9. Alexandria social media exploration process: Ingestion and initial analytics in the background; Domain Modeling using text analytics; a broad variety of Social Media Analytics; and identification of actionable insights through visualizations. Insights can lead to iterative modifications of the Domain Model and application of further analytics.

system [10], a precursor to IBM’s Social Media Analytics product [22], but other systems could be used). The results are placed into a LUCENE SOLR main-memory index to enable rapid searching, including against the Tweet text bodies, a key enabler for goals **LG4**, **SG1**, **SG3**, and **SG5**. The other background processing is to ingest, process, and index various background corpora to support text analytics. As described in further detail in Section V below, this is used to support the interactive domain scoping activity, relevant to goals **LG2**, **SG1**, **SG2**. And as describe in Section VI, this is also used to support the anomalous topics analytics and view (goals **LG2**, **SG3**).

Referring again to Figure 9, once a Domain Model is established for a Project, the Social Media analytics processing is performed. This is described in more detail in section VI below. After extraction and annotation, the desired analytics are invoked through REST APIs by an orchestration layer and the results are again placed into CouchDB. Finally, these can be accessed through several interactive visualizations.

As illustrated in Figure 10, most capabilities in Alexandria are accessed through REST services, which is the basic approach to supporting goals **LG1**, **SG3** and **SG6**. For capabilities involving large data volumes, the data is passed “by reference” for increased performance (**LG4**, **SG5**). At present the REST services are grouped more-or-less according to the architectural flow of Figure 9. (It is planned to REST-enable the background processing.) The REST services rely on a shared logical Data Store, which is currently comprised of LUCENE and CouchDB. This can be extended to other storage and access technologies without impacting the REST interfaces (goals **LG1**, **LG4**, **SG5**).

The REST-based architecture has already been applied to enable a rapid integration of Alexandria capabilities with IBM Research’s SystemG [13], a graph-based system that also supports social media analytics. In particular, the Alexandria Domain Models are now accessible to SystemG services, and the SystemG UI has been extended to support both Domain Scoping and Alexandria analytics views.

Alexandria exists as a software layer that can access raw repositories and streams of social media (and other) data, and that resides on top of several application, middleware, and data storage technologies. The system currently uses the GNIP Twitter reader and Board reader to access social media and web-accessible data. The application stack is

currently based on LUCENE, CouchDB, and HDFS for data storage and access, Hadoop for cluster management, IBM’s Big Insights, SPSS, and Social Media Analysis for analytics, and finally TomCat and Node.js to provide application server middleware. Alexandria lives above these layers, and could be extended to take advantage of other server capabilities (goals **LG1**, **LG4**, **SG3**, **SG5**).

V. DOMAIN SCOPING

Domain Scoping addresses the challenge of constructing Domain Models. A Domain Model is typically represented as families of keywords and composite topics (a.k.a., text extractors), which get applied to the corpus of text documents to realize the search or filtering in the corpus. Traditionally, Domain Scoping is performed by a subject matter expert who understands the domain very well and can specify precisely what the particular queries and search criteria should be for a given set of topics of interest. A central goal of Alexandria is to simplify significantly the task of creation of Domain Models as well as to lower the required domain expertise of the person creating Domain Models. To achieve that, we developed several techniques that leverage text analysis and data mining in order to assist at discovery and definition of relevant topics that will drive creation of search queries. In particular, we describe our approach for (1) discovery of relevant collocated terms, for (2) term clustering, and for (3) similar term generation. As illustrated in Section III, these three techniques combined together allow very easy, iterative definition of terms and topics (i.e., sets of collocated terms) relevant for a particular domain with minimal input required from the user. Other scoping tools can be incorporated into Alexandria, e.g., a tool based on using an ontology such as DBPedia.

A. Collocated Term Discovery

Alexandria employs two techniques – term frequency-inverse document frequency (TF-IDF) score and collocation – to discover significant relevant terms to a specific set of seed terms. Simply put, what Alexandria does is find documents that seed terms appeared within. This is called the “foreground” documents. It then harvests other terms that were mentioned in the documents and computes their significance.

To support this analytic, we acquired sample documents – documents considered general and representative enough of many different topics and domains – as the “background” materials for this operation. For this purpose we collected a complete week of documents (Sept 1-7 2014) from BoardReader. This extraction amounts to about 9 millions documents. The documents were then indexed in SOLR [24], a fast indexing and querying engine based on Lucene, for later fast access. Next we queried “NY Times” from this large set of documents, which resulted in news articles in many different areas including politics, sports, science and technology, business, etc. This set of documents is used to build a dictionary of terms that are not limited to a specific domain within a small sample. It is the basis for Alexandria to calculate term frequency in general documents.

From the foreground materials, Alexandria computes the significance of other terms in the documents using TF-IDF scores. TF-IDF score is a numerical statistic widely used in information retrieval and text mining to indicate

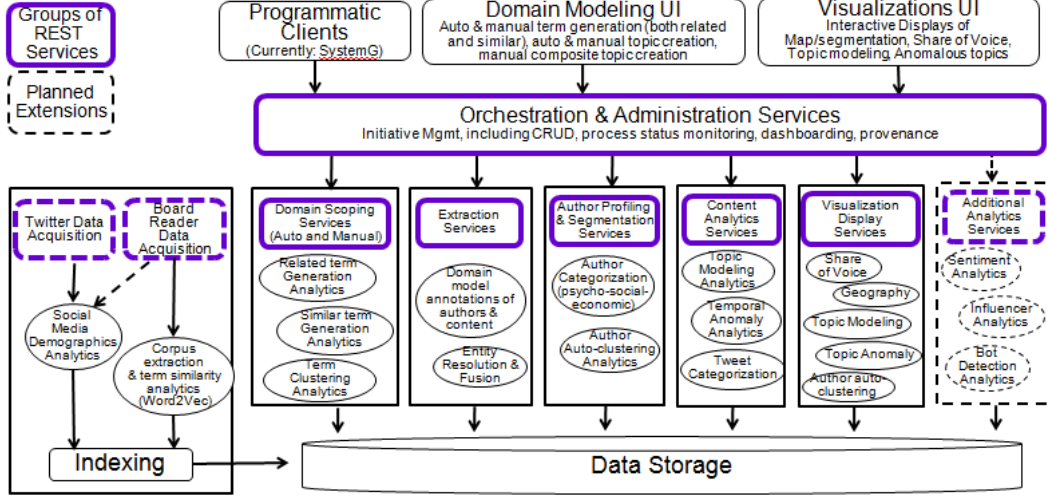


Figure 10. Alexandria supports loosely coupled RESTful services that orchestrate and invoke many functionalities, all sharing a common data store

the importance of a term to a document [14]. The score of a term is proportional to the frequency of the term in a document, but is offset by the frequency of the same term in general documents. The TF-IDF score of a word is high if the term has high frequency (in the given document) and a low frequency in the general documents. In other words, if a term appears a lot in a document, it may be worth special attention. However, if the term appears a lot in other documents as well, then its significance is low.

$$TF-IDF = TF(t, d) \times IDF(t, D)$$

$$IDF(t, D) = \log \frac{N}{|t \in D|}$$

A collocation is an expression consisting of two or more words that corresponds to some conventional way of saying things. They include noun phrases such as “weapon of mass destruction”, phrasal verbs like “make up” and other stock phrases such as “the rich and powerful.” We applied collocation to bring in highly relevant terms as phrases when the words collocate in the document and would make no sense as individual terms. More details of this technique can be found in [3]. Examples of these phrases are seen in Figure 1, for example, “small business,” “retail categories,” and “men shirts.”

For collocated term generation, the larger the corpus and the more accurate the results will be. However a very large corpus will suffer from efficiency and is not practical to use in an interactive environment such as Alexandria. Our hypothesis is that a week of general documents as a background corpus is a good enough representative of the bigger corpus, but is small enough to calculate the TF-IDF and collocation scores in a responsive manner.

B. Term Clustering and Similar Term Generation

Alexandria uses a term-clustering algorithm based on semantic similarities between terms to semantically group them into appropriate and strong “topics”. Alexandria uses Neural Network Language Models (NNLMs) that map words and bodies of text to latent vector spaces. Since they were initially proposed [3], a great amount of progress has been made in improving these models to capture many complex types of semantic and syntactic relationships [15],

[18]. NNLMs are generally trained in an unsupervised manner over a large corpus (greater than 1 billion words) that contains relevant information to downstream classification tasks. Popular classification methods to extract powerful vector spaces from these corpora rely on either maximizing the log-likelihood of a word, given its context words [15] or directly training from the probabilistic properties of word co-occurrences [18]. In Alexandria, we train our NNLMs on either a large corpus of Tweets from Twitter or a large corpus of news documents to reflect the linguistic differences in the target domain the end user is trying to explore. We also extended the basic NNLM architecture to include phrases that are longer than those directly trained in the corpus by introducing language compositionality into our model [23], [7], [16]. This way, our NNLM models can map any length of text into the same latent vector spaces for comparison.

The similarity measure obtained to support the term clustering is also used to generate new terms that are “similar” to the terms already in a topic.

VI. ANALYTICS VIEWS

This section briefly surveys two of the four main analytics algorithms currently supported by Alexandria; the others are omitted due lack of space.

A. Profile Extraction

As a pre-cursor to the other analytics in Alexandria, the tweets identified by the composite topics are extracted from the SOLR index and the corresponding authors’ profiles are compiled. Both the tweets and profiles are annotated along the composite topics, and stored for the Project in both CouchDB (noSQL database) and SOLR indexes. Alexandria incrementally fetches from the Twitter decahose to maintain a 6-month rolling window of tweets. We also incrementally perform analytics to compile authors’ user profiles. Attributes such as locations (used in showing geographic distribution), whether authors are parents, and intent to travel, are computed using tweets as evidence. The analytics based on previous research work done at IBM [10] has shown to show around 82% to 94% accuracy.

We provide a brief illustration of the running time of various steps. The current system is focused on a fixed set of

English-language Tweets from the Twitter Decahose (10% of all Tweets). With regards to background ingestion and initial processing, the current Alexandria infrastructure uses a 4 node cluster, with 1 as master and 3 as slaves; each node has 64MB of memory. We focus on the time needed to process through Alexandria. If a serialized machine were to be used then the extraction would be about 15 hours; With 10 nodes and 80 mappers there is a strong time reduction down to about 2 hours. Increasing to 17K mappers (the maximum number) brings the time to about 1 hour.

We also measured the end-to-end clock time for performing the extraction and annotation stage for a set of tweets. With a corpus of almost half a million tweets (452,201) the elapsed time was 4 minutes 29 seconds. With a corpus of almost a million tweets (949,241) it took 11 minutes and 31 seconds. (The numbers are not linear probably because the system is running on cloud-hosted virtual servers, which are subject to outside work loads at arbitrary times.) The processing includes writing the formatted data into both a CouchDB and a SOLR database. Looking forward, we expect to move towards an architecture with a single indexed data store, so that we can perform the annotations “in-place”.

B. Temporal Anomaly

Lastly, Alexandria performs topic analytics to help the user explore the topics discussed among tweets. Unlike many available topic detection algorithms [17], we define anomalous topics as terms that suddenly receive attention in a specific week when compared to the rest of the weeks in the data set. Alexandria uses a technique similar to the event detection domain [2]. It extracts terms from tweets, compute TF-IDF scores and frequencies and only retain terms with high TF-IDF score and high frequency. To calculate anomaly score for a term, we consider the frequency of the term in each week and its frequency over all the weeks in the data set. If the term’s frequency and score deviate a lot in a particular week from what it normally has over all, the term is considered anomalous. There could be an event or an emerging trend that caused the buzz, and hence people discuss more about the term in that week. This can trigger the user to look further to correlate research on events in that week. Following shows the formulas used for the calculation.

$$\begin{aligned} \text{anomalyScore}(term_i, week_j) &= \frac{\text{normFreq}(term_i, week_j)}{\text{normFreq}(term_i, \text{all_weeks})} \\ \text{normFreq}(term_i, week_j) &= \frac{\text{count}(term_i, week_j)}{\text{maxCount}(week_j)} \\ \text{normFreq}(term_i, \text{all_weeks}) &= \frac{\text{count}(term_i, \text{all_weeks})}{\text{maxCount}(\text{all_weeks})} \end{aligned}$$

VII. META-DATA SUPPORT FOR ITERATIVE EXPLORATION

Alexandria has been designed to support rapid, iterative, collaborative exploration of a domain including the usage of multiple analytics (goals **LG3**, **SG4**, **SG6**). This is enabled in part by the disciplined use of REST APIs to wrap the broad array of analytics capabilities (see Figure 10). But the fundamental enabler is the strongly data-centric approach taken for managing the several Projects that are typically created during the investigation of a subject area.

Data about all aspects of a Project (and pointers to more detailed information) is maintained in a CouchDB document, called *ProjectDoc*; this can be used to support a dashboard

about project status, and to enable invocation of various services. For example, the *ProjectDoc* holds a materialized copy of the domain model used to select the tweets and authors that are targeted by the Project. It maintains a record of which analytics have been invoked, and also maintains status during the analytics execution to enable a dashboard to show status and expected completion time to the end-user. Provenance data is also stored, to enable a determination of how data, analytics results, and visualizations were created in case something needs to be reconstructed or verified.

The *ProjectDoc* provides a foundation for managing flexible, ad hoc styles of iterative exploration. For example, with the *ProjectDoc* it is easy to support “cloning” of a Project to create a new one, and to combine the Topics and Composite Topics from multiple Projects to create a new one. It also allows for maintenance of information about whether analytics results have become out-of-date, and to support the incremental invocation of analytics, e.g., as new tweets become available. It also supports the inclusion of new Composite Topics into a Project’s domain model, along with controlled, incremental computation of the analytics for these additions.

VIII. RELATED WORK

Many papers focus on understanding social media. Various social media studies provide understanding of how information is gathered. For instance, [12] analyses community behaviors of social news site in the face of a disaster, [5] studies information sharing on Twitter during bird flu breakout, and [6] studies how people use search engines and twitter to gain insights on health information, providing motivation for ad hoc exploration of social data. Fundamentally, the authors of [19] elaborated on design features needed in a tool for data exploration and analysis, and coined the term “Information Building Applications.” They emphasized the support for tagging and categorizing raw data into categorization and the ability to restructure categories as their users, students, understand more about the data or discover new facts. The authors also emphasized the necessity of supporting fluid shift between concrete (raw data) and abstract (category of data) during the validation and iteration process, especially when faced with suspicious outcomes. While the paper discussed specifically about a tool for exploring streams of images, the nature of the approach is very similar to the process of exploring social media we are supporting in Alexandria. From another direction, as discussed in [8], an environment for analytics exploration, and application of the results, must support rich flexibility for pro-active knowledge-workers, and incorporate best practice approaches including Case Management and CRISP-DM [21] at a fundamental level. Because project management in Alexandria is based on data-centric principles (Section VII), along with the services-API-centric design, the system lays the foundation for the next generation of support for the overall analytics lifecycle.

Another novelty in our work is the combination of various text analytics and social media exploration tools into a broad-based solution for rapid and iterative domain modeling. While many tools exist, such as Topsy [25], Solr [24], Banana [1], we discovered that these tools do not support well the process and the human thoughts in gathering quality results. The existing tools typically tend

to aid in a fraction of the overall exploration task needed. More comprehensive, commercial tools such as HelpSocial [9] and IBM Social Media Analytics [22] are geared towards a complete solution. However, these tools require employing a team of consultants with deep domain expertise to operate as consulting services. Their support for the exploration process is not trivial and relies heavily on human labor and expertise. In terms of the research literature, Alexandria is helping to close a key gap in research on tooling for data exploration that was identified in [4].

IX. CONCLUSIONS AND DIRECTIONS

This paper describes the Alexandria system, which provides a combination of features aimed at enabling business analysts and subject matter experts to more easily explore and derive actionable insights from social media. The key novelties in the system are: (a) enabling iterative rapid domain scoping that takes advantage of several advanced text analytics tools, and (b) the development of a data-centric approach to support the overall lifecycle of flexible, iterative analytics exploration in the social media domain.

The Alexandria team is currently working on enhancements in several dimensions. Optimizations are underway, including a shift to SPARK for management and pre-processing of the background corpora that support the rapid domain scoping. Tools to enable comparisons between term generation strategies and other scoping tools are under development. A framework to enable “crowd-sourced” evaluation and feedback about the accuracy of extractors is planned. The team is working to support multiple kinds of documents (e.g., forums, customer reviews, and marketing content), for both background and foreground analytics. The team is also developing a persistent catalog for managing sets of topics and extractors; this will be structured using a family of industry-specific ontologies.

More fundamentally, a driving question is how to bring predictive analytics into the framework. A goal is to provide intuitive mechanisms to explore, view and compare the results of numerous configurations of typical machine learning algorithms (e.g., clustering, regression). This appears to be crucial for enabling business analysts (as opposed to data scientists) to quickly discover one-off and on-going insights that can be applied to improve business functions such as marketing, customer support, and product planning.

ACKNOWLEDGEMENTS

We would like to acknowledge other team members, Richard Goodwin, Sweefen Goh and Chitra Venkatramani. We also would like to acknowledge our colleagues from the SystemG project [13], including in particular Ching-Yung Lin, Danny Yeh, Jie Lu, Nan Cao, Jui-Hsin (Larry) Lai, and Roddrik Sabbah.

REFERENCES

- [1] Banana Development Team. Banana. <https://docs.lucidworks.com/display/SiLK/Banana>.
- [2] H. Becker, M. Naaman, and L. Gravano. Beyond trending topics: Real-world event identification on Twitter. In *Proc. Intl. Conf. on Web and Social Media (ICWSM)*, pages 438–441, 2011.
- [3] Y. Bengio et al. A neural probabilistic language model. *J. of Machine Learning Research*, 3:1137–1155, 2003.
- [4] E. Bertini and D. Lalanne. Investigating and reflecting on the integration of automatic data analysis and visualization in knowledge discovery. *SIGKDD Explorations*, 11(2):9–18, 2009.
- [5] C. Chew and G. Eysenbach. Pandemics in the age of Twitter: Content analysis of tweets during the 2009 H1N1 outbreak. *PLOS ONE*, 5(11), 2010. e14118. doi:10.1371/journal.pone.0014118.
- [6] M. Choudhury, M. Morris, and R. White. Seeking and sharing health information online: Comparing search engines and social media. In *Proc. ACM Intl. Conf. Computer Human Interaction (CHI)*, pages 1365–1375, 2014.
- [7] C. Goller and A. Kuchler. Learning task-dependent distributed structure-representations by backpropagation through structure. *IEEE International Conference on Neural Networks*, pages 347–352, 1996.
- [8] F. F. Heath III and R. Hull. Analytics Process Management: A new challenge for the BPM Community. In *BPM Workshops and Doctoral Consortium 2014 (LNC5)*, 2014.
- [9] HelpSocial Development Team. HelpSocial. <https://www.helpsocial.com/>.
- [10] M. Hernandez et al. Constructing consumer profiles from social media data. In *Proc. IEEE Big Data Conf.*, pages 710–716, 2013.
- [11] KQED. Disneyland measles outbreak hits 59 cases and counting. <http://www.npr.org/blogs/health/2015/01/22/379072061/disneyland-measles-outbreak-hits-59-cases-and-counting>.
- [12] A. Leavitt and J. Clark. Upvoting Hurricane Sandy: Event-based new production processes on a social news site. In *Proc. ACM Intl. Conf. Computer Human Interaction (CHI)*, pages 1495–1504, 2014.
- [13] C.-Y. Lin et al. Ibm system g home page. <http://systemg.research.ibm.com/index.html>.
- [14] C. Manning and H. Schütze. *Foundations of Statistical Natural Language Processing*. MIT Press, Cambridge, MA, 1999.
- [15] T. Mikolov, K. Chen, G. Corrado, and J. Dean. Efficient estimation of word representations in vector space. *CoRR*, abs/1301.3781, 2013.
- [16] T. Mikolov, M. Karafiát, L. Burget, J. Cernocký, and S. Khudanpur. Recurrent neural network based language model. In *Proc. 11th Ann. Conf. of the Intl. Speech Communication Association (INTERSPEECH)*, pages 1045–1048, September 26–30 2010.
- [17] M. Naaman, H. Becker, and L. Gravano. Hip and trendy: Characterizing emerging trends on Twitter. *J. Amer. Society for Inf. Science and Tech.*, 62(5):902–918, 2011.
- [18] J. Pennington, R. Socher, and C. D. Manning. Glove: Global vectors for word representation. In A. Moschitti, B. Pang, and W. Daelemans, editors, *Proc. 2014 Conf. Empirical Methods in Natural Language Processing, EMNLP 2014, Doha, Qatar*, pages 1532–1543. ACL, 2014.
- [19] A. Rajaraman and J. D. Ullman. *Mining of massive datasets*. Cambridge University Press, 2011.
- [20] A. Schuchat. Centers for disease control and prevention: Anne schuchat. <http://www.cdc.gov/about/leadership/leaders/schuchat.htm>.
- [21] C. Shearer. The CRISP-DM model: The new blueprint for data mining. *J. Data Warehousing*, 5(4):13–22, 2000.
- [22] SMA Development Team. IBM Social Media Analytics. <http://www-01.ibm.com/software/analytics/solutions/customer-analytics/social-media-analytics/>.
- [23] R. Socher et al. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the conference on empirical methods in natural language processing (EMNLP)*, volume 1631, page 1642. Citeseer, 2013.
- [24] SOLR Development Team. SOLR Home Page. <http://lucene.apache.org/solr/>.
- [25] Topsy Development Team. Topsy. <http://topsy.com>.